Army Aviation Fusion of Sensor-Pushed and Agent-Pulled Information

Angela M. Pawlowski and Craig W. Stoneking

Lockheed Martin Advanced Technology Laboratories
Camden, NJ 08012
[apawlows, cstoneki]@atl.lmco.com

Abstract

Supporting the Aviation Applied Technology Directorate's (AATD) Rotorcraft Pilot's Associate Advanced Technology Demonstration Program, the Lockheed Martin Advanced Technology Laboratories (ATL) demonstrated that formation of a Common Tactical Picture (CTP) from onboard and offboard sensors, via data fusion, was essential to automated decision aiding. In this case, the CTP was formed solely by the fusion of data provided by sensors dedicated to ownship tasks and from offboard sensors whose data was received via onboard processing systems, like JTIDS and the Improved Data Modem. Future network connectivity of aviation decision aiding systems to Army Battle Command System elements poses an opportunity to significantly enhance sensor data-only CTPs, if capability to autonomously and persistently discover and retrieve information from these stovepipe systems is applied. Addressing this challenge, the Lockheed Martin Advanced Technology Laboratories (ATL) has leveraged two autonomy-enabling technologies – multi-sensor data fusion and mobile intelligent agents, for Army aviation fusion of sensor-pushed and agent-pulled information. Over $14M in contracts and ATL research and development was combined to demonstrate these technologies in an Army ACT II proof-of-concept demonstration at the Air Maneuver Battle Lab. ATL plans to extend this concept for the AATD Airborne Manned and Unmanned System Technology Science and Technology Objective and the AATD Hunter Standoff Killer Team Advanced Concept Technology Demonstration.

Introduction

As the developers of the Rotorcraft Pilot's Associate (RPA) Data Fusion System, which was flight demonstrated aboard an AH/64-D in August 1999, we are keenly aware of the limitations of providing a Common Tactical Picture (CTP) for aviation decision aids that is based solely on the correlation and fusion of sensor data streamed to the aircraft. If the sensor data flowing into a tactical platform is inaccurate, time late, corrupted, or simply missing, the fused product of this data will necessarily be inaccurate, time late, or incomplete, as well. Given the wealth of information available within the Army Battle Command System (ABCS) of systems, and the emergence of our DARPA- and internally-funded Extendable Mobile Agent Architecture (EMAA), we have begun to exploit ABCS to enhance our production of the CTP.

Through a two-year internal research and development program, the Lockheed Martin Advanced Technology Laboratories (ATL) integrated its RPA Data Fusion software with its EMAA intelligent, mobile agents to bring ABCS data into the fusion process. Agents are small, mobile, autonomous software packages that travel over low-bandwidth, wireless, or conventional networks to remote hosts where they can deliver, collect, and analyze information, send reports, and issue new queries. This research and development program benefits from our previous 14 DARPA and DoD Intelligent Agent programs and over $1M of internal funding.

A CTP created from only incoming sensor data on Army aviation platforms is limited by the capabilities of the sensors, the reporting platform's intervisibility to battlefield entities, and the tactical maneuvers of the reporting platform. Frequent masking and unmasking, and shutting off sensors to maintain stealth, contribute to incomplete battlefield assessments. A wealth of pertinent intelligence exists within the ABCS, but as stovepipes designed for specific users, this information is rarely if ever exported to Army aviation. In cases where network connectivity exists, the challenge is to provide a means of monitoring these sources for pertinent information, and for selectively identifying and retrieving information within bandwidth tolerances.

Three challenges were addressed:
1) How to determine the outstanding needs of the CTP
2) How to discover information that met those needs

Presented at the American Helicopter Society 57th Annual Forum, Washington, DC, May 9-11, 2001.

in the Army's Joint Common Data Base (JCDB) with agents.

3) How to fuse the agent-gathered information into the CTP.

First, we developed an intelligent function that assessed the fused trackfile output from Data Fusion for clusters of tracks that contained either highly inaccurate positions or little to no classification information. Then, those clusters were translated into geographic regions for which mobile agents were deployed to search the JCDB, the Army's planned repository for information from ABCS systems, for either more accurate or better classification information. In addition, geographic regions of high interest were specified for agent search in JCDB.

The second challenge was met by creating itineraries, or directives, for mobile agents such that they could expeditiously travel over a network to JCDB and provide sentinel behavior, characterized by persistent queries to this constantly updated data source. As relevant track information was discovered by sentinel agents, it was reported back to the Data Fusion system, for incorporation into the CTP.

The third challenge was met by converting the JCDB information to meet Data Fusion input specifications. Since RPA Data Fusion was designed to receive real time sensor data from onboard and offboard sources, the agent-gathered information needed to be properly time-aligned before merging it with the CTP tracks. A direct application of this intelligent agent-enhanced Data Fusion process was delivered to the Air Maneuver Battlelab by Lockheed Martin Federal Systems – Owego and ATL as a part of the 2000 Army ACT II Battle Commander's Decision Aid.

A brief discussion of ATL's Data Fusion and Mobile Intelligent Agents, and the challenges addressed in their combination follows.

What is Data Fusion?

The task of battlefield situation assessment requires the ability to take reports from a variety of sensors, (RADAR, Infrared, IFF, etc.), and combine them into a single composite view of the position, movement, and identification of all of the targets, (tanks, ships, aircraft, etc.), within the area under observation. To illustrate, consider the situation depicted in Figure 1. Figure 1 presents a hypothetical battlefield situation display onto which reports from two different sensors, S1 and S2, are plotted. Due to errors inherent in the sensor measurements, the plotted position of each sensor report really only represents the center of an ellipse that defines a region in which the actual position of the target is expected to lie, with some high confidence. Given that the actual positions of the targets may be significantly displaced from the sensor-reported positions, and that some targets may be invisible to some sensors, it may be reasonable to interpret the sensor reports to represent two, three, or four actual targets on the battlefield. The problem of sensor data fusion is to choose the best interpretation of the collection of available sensor reports.

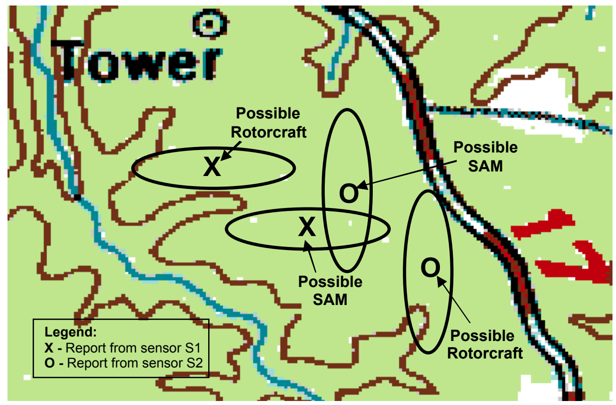

Figure 1. *Battlefield situation display with plots, (including error ellipses), from two sensors.*

In the case of our example it is readily apparent, from the relationships between the error ellipses, that the best interpretation of the scene is likely to be as depicted in Figure 2. Notice that the two center-most sensor reports, or tracks, were interpreted to represent a single actual target. By comparing the expected-error regions of the two sensor tracks, the expected error in the position of the resultant "fused" track has been greatly reduced. Notice also that the corroboration between the two sensors, in the classification of the target, results in a higher confidence in the classification of the fused track. The result of sensor data fusion is a single, de-cluttered representation of the battlefield with every known target plotted only once, with higher accuracy than could be achieved with any single sensor.

In such a simple case, it is not difficult for a human to interpret the scene mentally, without the aid of automation. However, as the numbers of sensors and targets increase to realistic values the complexity of the problem quickly increases to a level that requires automation. At ATL, we have developed an automated sensor data fusion engine with the demonstrated ability to fuse, in real time, data from as many as fourteen sensors, on as many as two hundred targets.

While automated sensor data fusion is a vast improvement over the simple union of all available sensor data, it still suffers from the same garbage-in-garbage-out limitation. If data on a given target is poor

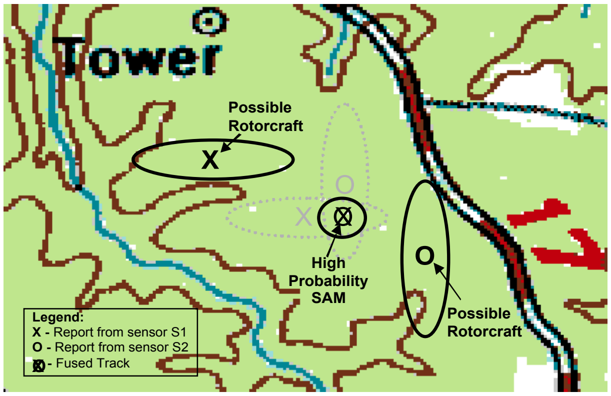

Figure 2. *The result of fusion is a de-cluttered display with higher confidence in position and identification of targets on the battlefield.*

or missing in the data from all available sensors, the data on that target is still likely to be poor or missing in the result of fusion. Our approach to this problem is to employ intelligent mobile agents to search for additional information that may be pulled from non-reporting data sources, (e.g., intelligence databases), to augment the fused sensor data and improve the tactical picture.

What is an Intelligent Mobile Agent?

An Agent is a software construct that is able to interact with its environment in order to perform tasks on behalf of the user, to further the user's goals. Mobility implies that the agent is able to travel between nodes of a network in order to make use of resources that are not locally available. Intelligence, in this context, implies that the agent is imbued with some degree of knowledge of its environment or subject domain that allows it to make decisions that affect its behavior, in response to changing characteristics of its environment, or the problem at hand.

At ATL we have developed the Extendable Mobile Agent Architecture, (EMAA), which provides an infrastructure for the deployment of intelligent mobile agents. The major components of EMAA are depicted in Figure 3.

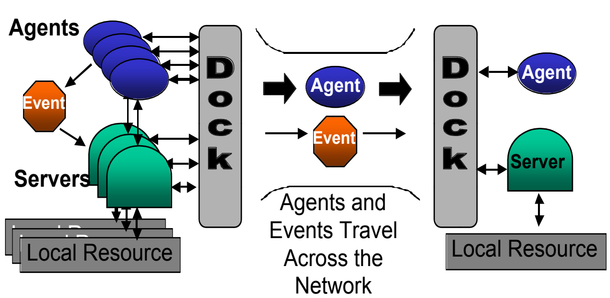

Figure 3. *Components of the Extendable Mobile Agent Architecture.*

A Dock is associated with a node of the network. It manages the transfer of agents to and from other Docks, and provides an execution environment in which agents can interact with Servers, or other Agents. Servers provide agents with services either directly (temporary data storage, specialized computation, etc.), or by acting as an interface to resources external to the Dock, (databases, web access, email, etc). Agents may communicate with Servers and other Agents by raising Events. These Events may either be local, in the style of the Java event model, or they may be Remote Events, which can travel between Docks. Agent behavior is controlled by an itinerary, which can be modeled as a sort of state-transition diagram in which each state specifies a task to perform, the node at which it should be performed, and some logic to determine what the next state should be.

Over the course of numerous DARPA and DoD applications, certain common patterns of agent behavior have emerged, and been abstracted for easy reuse. One such pattern is known as Sentinel Behavior. Sentinel Behavior is characterized by the persistent monitoring of a system for changes, then examining these changes to determine if they should be reported and, if so, reporting the new information to subscribed parties. This behavior may be implemented within the itinerary of a single agent, or it may be effected through the collaboration of multiple agents distributed throughout a widespread system.

Sentinel Behavior can be applied to the problem of improving the CTP with data from non-reporting data sources. An example of such an application is illustrated in Figure 4. In this system, the output of Sensor Data Fusion is collected to form the basis of the CTP. An Analysis Function examines the CTP to determine what additional information is needed. The Analysis Function dispatches an Investigation Agent to search for the needed information in a remote data source, such as JCDB. The results of the investigation are reported to the Fusion Input Interface, which feeds the new information into Sensor Data Fusion, as if it came from another sensor. The new information is then fused by Sensor Data Fusion into a new and improved CTP. This approach has been investigated in an internal research and development program and integrated into the ACT II Battle Commander's Decision Aid at the Air Maneuver Battle Lab.

Determining the Needs of the CTP

In Figure 4, the fused track data output from Sensor Data Fusion is collected in a WhiteBoard data structure on the Dock to form the basis of the Common Tactical Picture (CTP). The role of the Analysis Function is to examine the content of the CTP and identify regions

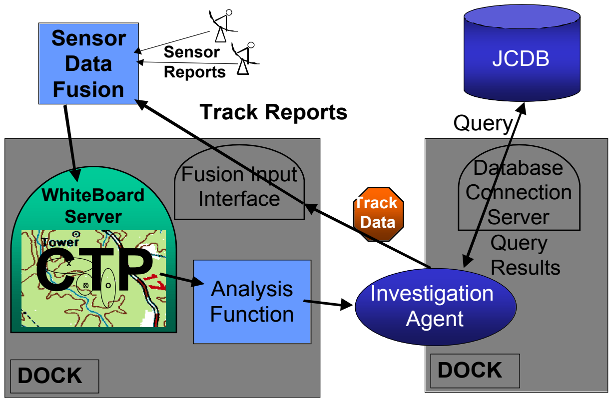

Figure 4. *An example of how agents, exhibiting Sentinel Behavior, can be used to improve the CTP with data from non-reporting data sources.*

where the CTP might be improved through special investigation. Types of information that we have exploited for identifying regions of interest include; pre-mission expectations, existing tracks with high error in position, existing tracks with missing classification information, and areas of suspicious activity.

The simplest of these is the use of pre-mission expectations. A region of the battlefield, along with a time, may be specifically called out by the user as a time and place of interest. The driving rationale for this is that a commander may have pre-battle expectations of enemy activity taking place in a particular area, at a particular time. The Analysis Function may be directed to select this region for investigation, at the given time, regardless of existing sensor tracks in the area, or it may be directed to investigate only if sensors have not detected the expected activity by the expected time.

A more complex strategy involves examining the existing fused tracks to identify those that may be improved by fusing with data from alternative data sources. To refine our agent's search spaces we adapted the linear-time clustering algorithm from our Sensor Data Fusion system to partition the existing trackfile into geographically defined clusters. Each cluster is then examined to see if the tracks within it are, collectively, of particularly poor quality. If the tracks within the cluster are deemed to have high positional error, or if class identification information is largely missing or of low certainty, then the cluster becomes a subject for investigation. A region that bounds the geographical extents of the subject cluster then becomes the region of interest for the investigation.

Exploiting suspicious behavior involves using the behavior of known tracks to focus the search for previously undetected tracks. For example, some sensors only detect a target while it is in motion, or only while it is transmitting radio signals. Consequently, by initiating or discontinuing either of these activities, the corresponding track may suddenly appear or disappear from the CTP. In the case of movement, it is not unreasonable to suspect that the target in question is arriving at or departing from some point of particular interest. For instance, such a point might be a fuel depot, a command center, or a congregation of other motionless, non-transmitting vehicles. Such points of interest, while not detectable by contributing sensors, may have been detected by other means of intelligence, so that information about them is available from some non-reporting intelligence data source. When a track suddenly appears or disappears in the CTP, the Analysis Function selects a region surrounding the point of appearance or disappearance as a region of interest to be investigated.

Information Discovery in the JCDB

Once the Analysis Function identifies a region of interest to investigate, it creates a new Investigation Agent. The Investigation Agent is given the latitude and longitude of two points that are opposite corners of a rectangle that defines the region of interest. Once set in motion, the Investigation Agent travels to a host Dock that provides access to the data source that is to be searched for information on tracks within the region of interest.

As our data source, we have used an instance of the Army's Joint Common Database, (JCDB). As the recipient of data from multiple ABCS systems, the JCDB contains a wealth of information relevant to the current Common Tactical Picture. For our AMBL ACT II demonstrations, we were concerned primarily with extracting JCDB information pertaining to the locations of enemy tracks. After obtaining access to the JCDB, the Investigation Agent forms a query to extract information on any enemy tracks located within the region of interest. The Investigation Agent's itinerary may be constructed such that it queries the JCDB only once, reports the results, and then either moves on, or simply dies. More commonly, the Investigation Agent exhibits true Sentinel Behavior by persisting (periodically repeating the query) to monitor for any updates to the information of interest. As the Investigation Agent discovers information of interest, it reports the information, via a remote event, to the Fusion Input Interface. The Investigation Agent keeps track of what information it has previously reported, so that it can be sure to report only new information from the results of subsequent queries to the JCDB.

Fusion of Sensor Pushed and Agent Pulled Data

Upon receipt of track data discovered by the Investigation Agent, the Fusion Input Interface feeds the data to Sensor Data Fusion. To do this, the Fusion Input Interface first preprocesses the data to translate position, velocity, error, and time data into the appropriate coordinate system and units for use by Sensor Data Fusion. The data that Sensor Data Fusion receives from the Fusion Input Interface is treated like any other sensor report. If the timestamp of the report is in the past, Sensor Data Fusion projects the position and error ellipse forward to the current time, before fusing with other sensor reports. With the incorporation of the agent-discovered data from JCDB, the output from Sensor Data Fusion now forms the basis of a new and improved CTP. The improvements to the CTP take the form of better position estimation and better class identification of existing tracks, as well as the appearance of previously undetected tracks. This improved CTP is then available for analysis by the Analysis Function, thereby effecting the continuous, iterative quality improvement of the CTP.

The Path Ahead

As Army Aviation evolves its decision aiding capabilities for mobile command teams and teaming with autonomous air vehicles, the need to present a CTP from mobile intelligent agents and information fusion has emerged. In programs like Airborne Manned and Unmanned System Technology (AMUST STO) and the Hunter Standoff Killer Team Advanced Concept Technology Demonstration, platforms will not always be equipped with sensor suites whose fused data become the CTP. The mobile command vehicle, for instance, is networked to the Army's ATCCS systems, but may or may not receive sensor tracks directly. Data Fusion techniques, coupled with mobile intelligent agents, will be employed to first discover, then fuse track information from ATCCS systems for the Commander's Associate CTP. Agents performing as sentinels to one platform's CTP will tailor and disseminate CTP subsets to platforms without access to this information to enhance situation awareness for all Hunter Standoff Killer Team participants.

Conclusions

In both our research and development effort and the AMBL ACT II project, mobile intelligent agents positively impacted the quality of the fused CTP. In the research and development effort, the CTP was enhanced with both missing and higher-quality threat and friendly track information. In the ACT II project, the Battle Commander's Decision Aid received specific types of threat information from the fusion of agent-discovered tracks that were discovered from a simulated ASAS data source writing to JCDB. This combination of Data Fusion and Mobile Intelligent Agents promises to be the foundation of the CTP for decision aiding in Army aviation.

Acknowledgements

The authors gratefully acknowledge the Aviation Applied Technology Directorate's scenario and problem definition support for the various internal research and development programs that have lead to this paper and for the continuing opportunities to apply fusion and agent technologies.

References

1. Hofmann, M.O., McGovern, A., and Whitebread, K.R., "Mobile Agents on the Digital Battlefield," Proceedings of the Second International Conference on Autonomous Agents, Minneapolis/St. Paul, MN, 1998.

2. McGrath, S., Chacón, D., and Whitebread, K., "Intelligent Mobile Agents in the Military Domain," Proceedings of the Autonomous Agents 2000 Workshop on Agents in Industry, Barcelona, Spain, 2000.

3. Pawlowski, A.M. and Gerken, P.M., "Simulator, Workstation, and Data Fusion Components for Onboard/Offboard Multi-Target Multi-Sensor Fusion," Proceedings of the 17th Digital Avionics Systems Conference, Seattle, WA, 1998.